## AMENDMENTS TO THE CLAIMS

This listing of claims replaces all prior versions and listings of claims in the application:

Listing of Claims:

1-40.   (Cancelled)

41.   (New)  A method for managing a data storage system, comprising:

configuring a first cache to perform at least one of the operations of retrieving data from and storing data at a first range of logical addresses (LAs) in a storage device;

configuring a second cache to perform at least one of the operations of retrieving data from and storing data at a first part of the first range of LAs;

configuring one or more third caches to perform at least one of the operations of retrieving data from an storing data at a second range of LAs in the storage device, the first, second, and third caches being implemented in three separate physical units, comprising software, and adapted to function substantially independently of each other;

detecting an inability of the second cache to retrieve data from or store data at the first part of the first range of LAs; and

reconfiguring at least one of the one or more third caches to perform at least one of the operations of retrieving data from and storing data at the first part of the first range of LAs in response to the inability while continuing to perform at least one of the operations of retrieving data from and storing data at the second range of LAs.

42.    (New)  The method according to claim 41, and comprising configuring one or more interfaces to receive input/output (IO) requests from host processor directed to specified LAs and to direct all the IO requests to the caches which have been configured to perform at least one of the operations of retrieving data from and storing data at the specified LAs.

43.    (New)  The method according to claim 42,

wherein the one or more interfaces comprise a mapping between the first and the second and the one or more third caches and the first and second ranges of the LAs,

wherein the one or more interfaces are adapted to convert the IO requests to one or more requests and to direct the one or more requests to at least one of the first and the second and the one or more third caches in response to the mapping, and

wherein detecting the inability comprises generating a reconfigured mapping between the first and the one or more third caches and the first and second ranges of the LAs, and directing the one or more request to at least one of the first and the one or more third caches in response to the reconfigured mapping.

44.    (New)  The method according to claim 41, wherein reconfiguring the at least one of the one or more third caches comprises processing data in the at least one of the first cache and the one or more third caches so as to restore the first cache and the one or more third caches to a state of full data redundancy.

45.    (New)  The method according to claim 44, wherein processing the data comprises classifying data in the first cache into a plurality of data groups.

46.     (New)  The method according to claim 45, wherein one of the data groups comprises

dirty data, and wherein processing the data comprises storing the dirty data at the one or more

third caches.


47.     (New)  The method according to claim 45, wherein one of the data groups comprises

dirty data, and wherein processing the data comprises storing the dirty data at the storage

device.


48.     (New)  The method according to claim 41, wherein reconfiguring the at least one of the

one or more third caches comprises retaining an initial configuration of the first cache.


49.     (New)  The method according to claim 41, wherein reconfiguring the at least one of the

one or more third caches comprises implementing a minimum redistribution of the first and the

second ranges among the first cache and the one or more third caches.


50.     (New)  The method according to claim 49, wherein implementing the minimum

redistribution comprises redistributing the first and the second ranges using a consistent hashing

function.


51.     (New)  The method according to claim 49, wherein implementing the minimum

redistribution comprises redistributing the first and the second ranges using a random number

function.

52.     (New)  The method according to claim 41, and comprising providing a system manager

which is adapted to configure the first, second and one or more third caches, to detect the

inability, and to reconfigure the at least one of the first cache and the one or more third caches.


53.     (New)  The method according to claim 52, wherein providing the system manager

comprises incorporating one or more manager processing units into at least one of the storage

device, the first cache, the second cache, and the one or more third caches, and operating the

one or more manager processing units in a cooperative manner.

54.    (New)  A data storage system, comprising:

a storage device wherein data is stored at logical  addresses (LAs);

a first cache which is configured to perform at least one of the operations of retrieving

data from and storing data at a first range of LAs in the storage device;

a second cache which is configured to perform at least one of the operations of

retrieving data from and storing data at the first range of LAs;

one or more third caches which are configured to perform at least one of the operations

of retrieving data from and storing data at a second range of LAs in the storage device, the first,

second, and one or more third caches being implemented in three separate physical units,

comprising software, and adapted to function substantially independently of each other; and

a system manager which is adapted to detect an inability of the second cache to retrieve

data from or store data at the first part of the first range of LAs, and which reconfigures at least

one of the one or more third caches to perform at least one of the operations of retrieving data

from and storing data at the first range of LAs in response to the inability while continuing to

perform at least one of the operations of retrieving data from and storing data at the second

range of LAs.

55.    (New)  The storage system according to claim 54, and comprising one or more

interfaces which are configured to receive input/output (IO) requests from host processors

directed to specified LAs and to direct all the IO requests to the caches which have been

configured to perform at least one of the operations of retrieving data from and storing data at

the specified LAs.

56.    (New)  The storage system according to claim 55,

wherein the one or more interfaces comprise a mapping between the first and the second and the one or more third caches and the first and second ranges of the LAs,

wherein the one or more interfaces are adapted to convert the IO requests to one or more requests and to direct the one or more requests to at least one of the first and the second and the one or more third caches in response to the mapping, and

wherein detecting the inability comprises the system manager generating a reconfigured mapping between the first and the one or more third caches and the first and second ranges of the LAs, and directing the one or more requests to at least one of the first and the one or more third caches in response to the reconfigured mapping.

57.    (New)  The storage system according to claim 54, wherein reconfiguring the at least one of the one or more third caches comprises processing data in at least one of the first cache and the one or more third caches so as to restore the first cache and the one or more third caches to a state of full data redundancy.

58.    (New)  The storage system according to claim 57, wherein processing the data comprises classifying data in the first cache into a plurality of data groups.

59.    (New)  The storage system according to claim 58,

wherein one of the data groups comprises dirty data, and

wherein processing the data comprises storing the dirty data at the one or more third caches.

60.　　(New)　The storage system according to claim 58,

wherein one of the data groups comprises dirty data, and

wherein processing the data comprises storing the dirty data at the storage device.

61.　　(New)　The storage system according to claim 54, wherein reconfiguring the at least one of the one or more third caches comprises the first cache retaining an initial configuration.

62.　　(New)　The storage system according to claim 54, wherein reconfiguring the at least one of the one or more third caches comprises the system manager implementing a minimum redistribution of the first and the second ranges among the first cache and the one or more third caches.

63.　　(New)　The storage system according to claim 62, wherein implementing the minimum redistribution comprises redistributing the first and the second ranges using a consistent hashing function.

64.　　(New)　The storage system according to claim 62, wherein implementing the minimum redistribution comprises redistributing the first and the second ranges using a random number function.

65.　　(New)　The storage system according to claim 64, wherein the system manager comprises one or more manager processing units which are incorporated into at least one of the storage device, the first cache, the second cache, and the one or more third caches, and wherein the one or more manager processing units operate in a cooperative manner.

66.     (New)  A storage system, comprising:

one or more mass storage devices, coupled to store data at respective first ranges of logical addresses (LAs);

a plurality of interim fast-access-time caches, each of the plurality of interim fast-access-time caches comprising software and configured to operate substantially independently of one another, each interim fast-access-time cache being assigned a respective second range of the LAs and coupled to receive data from and provide data to the one or more mass storage devices having LAs within the respective second range;

one or more interfaces, which are adapted to receive input/output (IO) requests from host processors directed to specified LAs and to direct all the IO requests to the interim fast-access-time cache to which the specified LAs are assigned; and

a further plurality of interim fast-access-time caches adapted to be configured to be assigned the respective second range of the LAs and coupled to receive data from and provide data to the one or more mass storage devices having LAs within the respective second range when any interim fast-access-time cache fails.


67.     (New)  The storage system according to claim 66, wherein the mass storage devices comprise one or more disks.

68.    (New)  A method for storing data, comprising:

storing the data in one or more mass storage devices having respective first ranges of

logical addresses (LAs);

assigning to each of a plurality of interim fast-access-time caches a respective second

range of the Las, each of the plurality of interim fast-access-time caches comprising software

and configured to operate substantially independently of one another; coupling the plurality of

interim fast-access-time caches to receive data from and provide data to the one or more mass

storage devices having LAs within the respective second range;

receiving input/output (IO) requests from host processors directed to specified LAs; and

directing all the IO requests to the interim fast-access-time cache to which the specified

LAs are assigned,

wherein a further interim fast-access-time cache is adapted to be configured to

be assigned the respective second range of the LAs and coupled to receive data from and

provide data to the one or more mass storage devices having LAs within the respective second

range when any interim fast-access-time cache fails.


69.    (New)  The method according to claim 68, wherein the mass storage devices comprise

one or more disks.

70.    (New)  A system for transferring data to and from one or more mass storage devices which store data at respective first ranges of logical addresses (LAs), comprising:

a plurality of interim fast-access-time caches, each of the plurality of interim fast-access-time caches comprising software and configured to operate substantially independently of one another, each interim fast-access-time cache being assigned a respective second range of the LAs and coupled to receive data from and provide data to the one or more mass storage devices within the respective second range;

one or more interfaces, which are adapted to receive input/output (IO) requests from host processors directed to specified LAs and to direct all the IO requests to the interim fast-access-time cache to which the specified LAs are assigned; and

a further plurality of interim fast-access-time caches adapted to be configured to be assigned the respective second range of the LAs and coupled to receive data from and provide data to the one or more mass storage devices having LAs within the respective second range when any interim fast-access-time cache fails.

71.    (New)  The system according to claim 70, wherein the mass storage devices comprise one or more disks.

72.    (New)  A method for transferring data to and from one or more mass storage devices which store data at respective first ranges of logical addresses (LAs), comprising:

assigning to a plurality of interim fast-access-time caches, each of the fast-access-time caches comprising software and configured to operate substantially independently of one another, respective second ranges of the LAs;

coupling the plurality of interim fast-access-time caches to receive data from and provide data to the one or more mass storage nodes having LAs within the respective second ranges;

receiving input/output (IO) requests from host processors directed to specified LAs; and

directing all the IO requests to the interim fast-access-time cache to which the specified LAs are assigned,

wherein a further interim fast-access-time cache is adapted to be configured to be assigned the respective second range of the LAs and coupled to receive data from and provide data to the one or more mass storage devices having LAs within the respective second range when any interim fast-access-time cache fails.


73.    (New)  The method according to claim 72, wherein the mass storage devices comprise one or more disks.